

NAVAL POSTGRADUATE SCHOOL

Monterey, California



THESIS

APPLICATION OF A SYSTEM-BASED INVENTORY MODEL TO MARINE CORPS REPAIRABLE PARTS

by

Craig P. Barnett

September 2001

Thesis Advisor:
Second Reader:


Kevin R. Gue
David A. Schradly

Approved for public release; distribution is unlimited.

Report Documentation Page

Report Date 30 Sep 2001	Report Type N/A	Dates Covered (from... to) -
Title and Subtitle Application of a System-Based Inventory Model to Marine Corps Repairable Parts	Contract Number	
	Grant Number	
	Program Element Number	
Author(s) Barnett	Project Number	
	Task Number	
	Work Unit Number	
Performing Organization Name(s) and Address(es) Research Office Naval Postgraduate School Monterey, Ca 93943-5138	Performing Organization Report Number	
Sponsoring/Monitoring Agency Name(s) and Address(es)	Sponsor/Monitor's Acronym(s)	
	Sponsor/Monitor's Report Number(s)	
Distribution/Availability Statement Approved for public release, distribution unlimited		
Supplementary Notes		
Abstract		
Subject Terms		
Report Classification unclassified	Classification of this page unclassified	
Classification of Abstract unclassified	Limitation of Abstract UU	
Number of Pages 52		

***Amateurs discuss strategy,
Professionals study logistics***



REPORT DOCUMENTATION PAGE			<i>Form Approved OMB No. 0704-0188</i>	
Public reporting burden for this collection of information is estimated to average 1 hour per response, including the time for reviewing instruction, searching existing data sources, gathering and maintaining the data needed, and completing and reviewing the collection of information. Send comments regarding this burden estimate or any other aspect of this collection of information, including suggestions for reducing this burden, to Washington headquarters Services, Directorate for Information Operations and Reports, 1215 Jefferson Davis Highway, Suite 1204, Arlington, VA 22202-4302, and to the Office of Management and Budget, Paperwork Reduction Project (0704-0188) Washington DC 20503.				
1. AGENCY USE ONLY (Leave blank)		2. REPORT DATE September-2001	3. REPORT TYPE AND DATES COVERED Master's Thesis	
4. TITLE AND SUBTITLE: Title (Mix case letters) Application of a System-Based Inventory Model to Marine Corps Repairable Parts			5. FUNDING NUMBERS	
6. AUTHOR(S) Barnett, Craig P.				
7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES) Naval Postgraduate School Monterey, CA 93943-5000			8. PERFORMING ORGANIZATION REPORT NUMBER	
9. SPONSORING / MONITORING AGENCY NAME(S) AND ADDRESS(ES) N/A			10. SPONSORING / MONITORING AGENCY REPORT NUMBER	
11. SUPPLEMENTARY NOTES The views expressed in this thesis are those of the author and do not reflect the official policy or position of the Department of Defense or the U.S. Government.				
12a. DISTRIBUTION / AVAILABILITY STATEMENT Approved for public release; distribution is unlimited.			12b. DISTRIBUTION CODE	
13. ABSTRACT (maximum 200 words) A critical component of the Marine Corps' self-sustainment capability is its ability to procure and repair components for its ground equipment fleets. Secondary repairables consist of components that can be repaired, and for which repair is generally more economical and timely than purchase. The Marine Corps currently maintains spare repairable parts at seven principal locations, each operating independently of the other. There is excess inventory Service-wide because of the isolation of the inventories and because of mathematical flaws in the Marine Corps' sparing methodology. The Marine Corps is seeking to centralize the management of secondary repairables and is considering options that include centralizing responsibility and funding (while keeping the inventory model as it is) and changing the inventory model as well as the responsibility and funding. We demonstrate that a centralized, "enterprise-wide" model of the inventory is superior to a decentralized one. Measures of comparison are total inventory cost and end-item availability. We evaluate stock levels calculated by both the current model and a commercial application called VMetric™-XL. For a selected end-item, the current model produces stock levels totaling \$25.9M in inventory and achieves 89.1% availability. At the same level of availability, VMetric recommends stock levels totaling \$2.9M, a stunning 89% reduction in cost. We explain these results and suggest implications for Marine Corps logistics support.				
14. SUBJECT TERMS Inventory Management, Marine Corps Secondary Repairables, Logistics, Supply Chain Management, Readiness-Based Sparing, Modeling and Simulation			15. NUMBER OF PAGES 50	
			16. PRICE CODE	
17. SECURITY CLASSIFICATION OF REPORT Unclassified	18. SECURITY CLASSIFICATION OF THIS PAGE Unclassified	19. SECURITY CLASSIFICATION OF ABSTRACT Unclassified	20. LIMITATION OF ABSTRACT UL	

THIS PAGE INTENTIONALLY LEFT BLANK

Approved for public release; distribution is unlimited

**APPLICATION OF A SYSTEM-BASED INVENTORY MODEL TO MARINE
CORPS REPAIRABLE PARTS**

Craig P. Barnett
Captain, United States Marine Corps
B.S., United States Naval Academy, 1992

Submitted in partial fulfillment of the
requirements for the degree of

MASTER OF SCIENCE IN OPERATIONS RESEARCH

from the


**NAVAL POSTGRADUATE SCHOOL
September 2001**

Author:


Craig P. Barnett

Approved by:


Kevin Gue, Thesis Advisor


David Schrady, Second Reader


James Eagle, Chairman
Department of Operations Research

THIS PAGE INTENTIONALLY LEFT BLANK

ABSTRACT

A critical component of the Marine Corps' self-sustainment capability is its ability to procure and repair components for its ground equipment fleets. Secondary repairables consist of components that can be repaired, and for which repair is generally more economical and timely than purchase. The Marine Corps currently maintains spare repairable parts at seven principal locations, each operating independently of the other. There is excess inventory Service-wide because of the isolation of the inventories and because of mathematical flaws in the Marine Corps' sparing methodology.

The Marine Corps is seeking to centralize the management of secondary repairables and is considering options that include centralizing responsibility and funding (while keeping the inventory model as it is) and changing the inventory model as well as the responsibility and funding. We demonstrate that a centralized, "enterprise-wide" model of the inventory is superior to a decentralized one. Measures of comparison are total inventory cost and end-item availability. We evaluate stock levels calculated by both the current model and a commercial application called VMetric™-XL.

For a selected end-item, the current model produces stock levels totaling \$25.9M in inventory and achieves 89.1% availability. For the same level of availability, VMetric recommends stock levels totaling \$2.9M, a stunning 89% reduction in cost. We explain these results and suggest implications for Marine Corps logistics support.

THIS PAGE INTENTIONALLY LEFT BLANK

TABLE OF CONTENTS

I.	INTRODUCTION.....	1
A.	BACKGROUND	1
B.	RESEARCH FOCUS.....	3
II.	THE MODELS.....	7
A.	CURRENT MODEL.....	7
1.	The Process	7
2.	Intermediate Level.....	8
3.	Depot	9
4.	The Model.....	10
B.	SYSTEM-BASED METHOD	11
1.	Single-Site Model	12
2.	Multi-Echelon Model.....	14
3.	Modifications and Generalizations.....	16
III.	DATA AND ASSUMPTIONS.....	17
A.	CURRENT MODEL.....	17
B.	DATA FOR VMETRIC™-XL.....	19
IV.	RESULTS AND ANALYSIS	21
A.	CHANGES WITHIN THE CURRENT MODEL.....	21
B.	CHANGING THE MODEL.....	23
C.	CONCLUSIONS	26
	LIST OF REFERENCES	29
	INITIAL DISTRIBUTION LIST	31

THIS PAGE INTENTIONALLY LEFT BLANK

LIST OF FIGURES

Figure 1.	Current Process Schematic.	7
Figure 2.	Sample Availability vs. Cost Curve.....	13
Figure 3.	Optimal Availability vs. Cost Curve for LAV Repairables	24

THIS PAGE INTENTIONALLY LEFT BLANK

LIST OF TABLES

Table 1.	Effect of Various Median Calculations.....	22
Table 2.	Comparison of Revised and Current Formulas.....	23
Table 3.	Direct Comparison of Current Model vs. VMetric.....	25
Table 4.	Modifying Expected Demands in the Current Model.....	26

THIS PAGE INTENTIONALLY LEFT BLANK

LIST OF ACRONYMS

CEC	Combat Essentiality Code
DLR	Depot Level Repairable
FLR	Field Level Repairable
ILC	Integrated Logistics Capability
IMA	Intermediate Maintenance Activity
LAR	Light Armored Reconnaissance (Battalion)
LAV	Light Armored Vehicle
LX	Analysis office within the Installations and Logistics Department of Headquarters, Marine Corps
MCLB	Marine Corps Logistics Base
MEF	Marine Expeditionary Force
METRIC	Multi-Echelon Technique for Recoverable Item Control
MFR	Maintenance Failure Rate
MIMMS	Marine Corps Integrated Maintenance Management System
MRR6	Mean Replacement Rate (failures per million operating hours)
MTBF	Mean Time Between Failures
OST	Order and Ship Time
PLT	Procurement Lead Time
PTRF	Peacetime Technical Replacement Factor
RBS	Readiness Based Sparing
RCT	Repair Cycle Time
RIP	Repairable Issue Point
RO	Requisitioning Objective
RR	Repair Rate
SASSY	Supported Activity Supply System
SMRC	Source, Maintenance and Recoverability Code
SRMO	Secondary Repairables Management Office
VARI-METRIC	variant of METRIC incorporating variance of demand

THIS PAGE INTENTIONALLY LEFT BLANK

ACKNOWLEDGMENTS

The author acknowledges and thanks the following people for their part in the successful completion of this thesis:

Jesus Christ, the source of all wisdom, knowledge and truth.

His wife, Janet, for her love and support, and for doing more than her share to keep the household running smoothly. This truly would not have been possible without her.

Professor Kevin Gue, for his advice and support, and for help through the rough spots.

Chris Hampson, Bob Butler and Jon Redfield from Systems Exchange, for advice, incredible technical support, and for the use of VMetric-XL.

THIS PAGE INTENTIONALLY LEFT BLANK


EXECUTIVE SUMMARY

A critical component of the Marine Corps' self-sustainment capability is its ability to procure and repair components for its ground equipment fleets. Secondary repairables consist of components that can be repaired, and for which repair is generally more economical and timely than purchase. In addition to those secondary repairables installed in equipment, the Marine Corps currently maintains allowances of \$500M in spare repairable parts at seven principal locations, each operating independently of the other. Each location attempts to maintain sufficient spares to satisfy local demand and enough safety stock to guard against local variability in demand, order lead time, and repair cycle time. There is excess inventory Service-wide because of the isolation of the inventories and because of mathematical flaws in the Marine Corps' sparing methodology.

As a result of the Integrated Logistics Capability initiative, the Marine Corps is seeking to centralize the management of secondary repairables and is considering options that include centralizing responsibility and funding (while keeping the inventory model as it is) and changing the inventory model as well as the responsibility and funding. We demonstrate that a centralized, "enterprise-wide" model of the inventory is superior to a decentralized one. Measures of comparison are total inventory cost and end-item availability. The centralized inventory model we use is the VARI-METRIC algorithm. We evaluate stock levels calculated by both the current model and a commercial application of VARI-METRIC called VMetric™-XL.

We use the Light Armored Vehicle (LAV) as the weapon system on which to perform our research. We limit the parts sample to those designated mission critical. We simplify the locations by combining the Marine Corps Logistics Bases into one site and considering only the three active Repairable Issue Points. We use data from Marine Corps systems and highlight several gaps and inconsistencies in these data.

For a selected end-item, the current model produces stock levels totaling \$25.9M in inventory and achieves 89.1% availability. For the same level of availability, VMetric recommends stock levels totaling \$2.9M, a stunning 89% reduction in cost. These results are summarized in the following chart.



Our results show that there are huge potential benefits from managing the Marine Corps inventory as a whole rather than as a disjoint group of independent inventories. Centralized management of secondary repairables cannot mean simply consolidating ownership under one command or funding all RIPs through a single command. To get the maximum benefits, Marine Corps secondary repairables must be treated as one inventory separated among geographically separate sites, not as separate inventories. Using the same modeling assumptions, we expect that similar results could be achieved with any other item of equipment in the Marine Corps inventory.

I. INTRODUCTION

A. BACKGROUND

A critical component of the Marine Corps' self-sustainment capability is its ability to procure and repair components for its ground equipment fleets. Secondary repairables consist of components that can be repaired, and for which repair is generally more economical and timely than purchase. Secondary repairables include a wide variety of items, from radar, computer and radio circuit cards to large, special-purpose diesel engines. In addition to those secondary repairables installed in equipment, the Marine Corps currently maintains allowances of \$500M in spare repairable parts at seven principal locations. At any given time, these spares are either in a serviceable condition in inventory (possibly deployed), or in one of three other states: (1) on order to a source of supply, (2) undergoing corrective maintenance, or (3) in transit between inventories. This last state notwithstanding, the seven principal inventories largely operate independently from each other. Therefore, each location attempts to maintain sufficient spares to provide assets in exchange for local in-process inventory or supply pipeline stock, as well as safety stock to guard against local variability in demand, order lead time, and repair cycle time. This causes excess inventory Service-wide because the quantity of safety stock required to buffer against stock-outs throughout the Marine Corps inventory is less than the sum of the safety stock quantities needed to buffer against stock-out at each location. Further, local management of stocks can lead to other inefficiencies. As of July 2000, ground equipment managers held \$416M in serviceable on-hand assets, \$84M in pending procurement orders, \$34M in the repair cycle, and \$16M in transit, totaling \$550M – 10% above total authorized allowances [1]. In addition, a Naval Audit Report states that in 1996, computed stock allowances were \$234M higher than necessary [2].

In 1998, the Marine Corps began a ground-up re-engineering of its logistics processes called Integrated Logistics Capability (ILC). One of the proposals resulting from the ILC conferences was to consolidate the management of repairable parts under Materiel Command. The idea is that efficiencies can be gained by having a common

“overhead” and by taking an “enterprise-wide” view of demand and stock levels. The Marine Corps has mandated a change to a centralized management policy for secondary repairables [3, 4] with Marine Corps Logistics Bases (MCLB) as the lead agency [5]. This initiative is in keeping with the Department of Defense Logistics Strategic Plan, which has stated objectives of reducing worldwide inventories and the implementation of a “virtual” inventory control point structure within each service component. In the Plan, the definition of a virtual inventory control point is a management structure for multiple, geographically separate inventory control points under a single command.

The Secondary Repairables Management Office (SRMO) has outlined four possible courses of action: (1) Status quo in terms of inventory policy but with a new chain of command, (2) Status quo in terms of inventory policy but with a new chain of command and funding relationship, (3) A “virtual inventory” in terms of inventory policy (i.e. stock determination) as well as command relationships and funding, and (4) Outsource the entire secondary repairable inventory process to a third party logistics provider [6]. Options (1) and (2) consider centralization only of responsibility and funding while leaving the management decisions (e.g. stock levels, purchases, repair decisions) decentralized at the various inventory locations. Our research compares option (3) to (1) and (2). In other words, we explore the benefits of using a consolidated inventory management policy as compared to the current policy of independently operating inventory agencies. The problem is to determine whether an inventory centralized in stock policy provides sufficient benefits to justify the additional effort to make the necessary changes. We demonstrate that a centralized inventory management process, specifically in terms of calculating stock levels, is superior to the current decentralized process.

A simple mathematical example illustrates the advantage of consolidating inventory. Consider three inventory sites, each of which experiences demands on a single item. Assume that the daily demands at each site look like random draws from a Poisson distribution with mean equal to five. If we want to hold stock such that we have only a five percent chance of stock out at each location, we would hold nine at each site for a total of 27 units. To consolidate the inventory, we consider the demand in aggregate. The sum of Poisson distributions is also a Poisson distribution with a mean

equal to the sum of the individual means. Therefore, the consolidated inventory would experience Poisson demands with a mean of 15 units. To achieve a five percent chance of stock-out in this case, we would need to hold only 22 units – five less than the distributed case.

B. RESEARCH FOCUS

A general concept of parts management that has gained a great deal of attention in recent time is readiness-based sparing (RBS). The idea is to stock those spares that provide the greatest contribution to the readiness of an end-item, or group of end-items. Instead of treating each part and each inventory location in isolation, RBS attempts to look across all inventory sites and all parts to increase the readiness of a weapon system. While there are many system-based inventory control methods, most are derived from or variations of the Multi-Echelon Technique for Recoverable Item Control (METRIC) model developed by Sherbrooke, originally published in 1968 [7]. We choose to use a METRIC-based model because such models have been proven in military and commercial applications over more than thirty years [8, 9]. We actually use a VARI-METRIC model [10, 11], which is an extension of METRIC that relaxes many of the restrictions and assumptions required for the basic METRIC model.

Using the same inputs, we calculate stock levels of the system components using both the current method and a system-based optimization model. We compare the results with regard to total inventory investment and end-item readiness. We also argue that a system-based model enables better decision-making than a piecemeal method.

For the purposes of this study, we chose to limit the scope to one end-item. This allowed us to focus on the readiness of that item, and the parts associated with it. We concentrate on the Light Armored Vehicle (LAV) because it is one in which the Marine Corps holds a great deal of interest and of which the Marine Corps is the primary user. Because of limitations in data availability, the scope expanded to almost all the variants of the LAV. It would have been impossible to segregate the part demands to a particular variant. To maintain commonality among the operating units, we only looked at the six most common LAV variants:

- LAV-25: Standard LAV with 25mm chain gun,

- LAV-AT: Anti-tank variant,
- LAV-C: Command and control variant,
- LAV-L: Logistics variant,
- LAV-M: Mortar variant, and
- LAV-R: Recovery and repair variant.

We ignore the electronic warfare and air defense variants because they are used only in special-purpose units. Also to limit the scope, the two depots (Albany and Barstow) are considered as one and the three Marine Expeditionary Forces are the only operational sites. Since we are primarily concerned with those components that affect the combat readiness of the end item, we limit the parts list to those with a combat essentiality code (CEC) of 5. We consider only operating stocks, not mount-out or war-reserve stocks, because the latter quantities are set by an entirely different process and are not used in daily operations.

The availability of accurate and complete data proved to be the major hurdle in this research. In particular, the Marine Corps does not retain operational usage data or any information connecting failures to usage (e.g. mean time between failure, MTBF). What data were available had to be collected from many, varied sources. Some data were not consistent. For example, a single part might have a different source, maintenance and recoverability code (SMRC) on each of several vehicle parts lists. Some data were very difficult to obtain. For example, part indenture information is kept primarily in hard-copy technical manuals and parts lists.

We use demand data from the period October 1997 to September 2000 extracted from the Marine Corps Integrated Maintenance Management System (MIMMS) to calculate the actual stock levels, also called requisitioning objectives (ROs), for each part in our sample at each location. We also calculate stock levels from these data using the revised version of the current formula that was derived by analysts in the LX branch at Headquarters, Marine Corps. The important elements for these calculations are the medians of monthly demands, monthly successful repairs, order-ship time (OST) and repair cycle time (RCT). From the ROs, we calculate the total inventory investment, the expected number of backorders, and from that, the supply availability for the whole

system (LAV) expected under the current and revised formulas. We use a spreadsheet to perform all calculations for the current and revised formulas.

For the centralized model, we calculate the stock levels using a commercial software package called VMetric™-XL. VMetric is built around the VARI-METRIC algorithm. We input, as much as possible, the same data elements used in the current formula. The program provides the stock level for each part at each location. It also calculates the optimal availability vs. cost curve for the total inventory and the availability. We then observe the difference between the investment levels at a common availability. Conversely, we also observe the difference between the availabilities at a given level of investment.

THIS PAGE INTENTIONALLY LEFT BLANK

II. THE MODELS

A. CURRENT MODEL

1. The Process

When a secondary repairable fails, the user takes the failed part to the local Repairable Issue Point (RIP) where a serviceable part is provided in one-for-one exchange, if one is available. If no serviceable items are available, a backorder is created. The failed part is sent to the appropriate repair facility. The local Intermediate Maintenance Activity (IMA) repairs some parts, some are sent directly to the Depot, and some are sent to the manufacturer. If an item cannot be repaired at the IMA, it is either disposed of (if it is beyond all repair), or it is sent to the Depot and a serviceable part is provided in direct exchange. If no serviceable part is available, a backorder at the Depot is created. When a failed part enters Depot repair it can either be successfully repaired and returned to Depot stock, or it can fail repair and be sent to disposal. If an item fails repair at the Depot, a new item is procured from the appropriate source. A schematic of this process is shown in Figure 1. Each level operates approximately on an $(s-1, s)$ inventory policy such that the quantity $OH + DI - BO$ is kept constant; where OH is the on-hand quantity, DI is the due-in quantity and BO is the backordered quantity.




Figure 1. Current Process Schematic.

2. Intermediate Level

At the RIPs, stock levels (called RO for requisition objective) are computed by the Marine Corps supply system called SASSY (Supported Activity Supply System) once per quarter using the following formula:

$$RO = RR \cdot RCT / 30 + (MFR - RR) \cdot OST / 30 + SL \quad (1)$$

$$\text{where, } SL = RR \cdot RCT / 30 + (MFR - RR) \cdot OST / 30 \quad (2)$$

and,

- SL: Safety Level,
- RR: Repair Rate – median number of items successfully repaired by the IMA each month,
- RCT: Repair Cycle Time – median number of days taken to repair an item ($30 \leq RCT \leq 90$),
- MFR: Maintenance Failure Rate – median monthly demands for exchange,
- OST: Order and Ship Time – median time to receive assets from supply source ($OST \geq 30$). [12]

A striking feature of this formula is how SASSY calculates the median demands (MFR). At the time of the computation, the system uses the previous 12 monthly observations of demand. The months with zero demand are removed from consideration. The median is calculated using only the months with positive demand. This artificially inflates demand because those months with zero demands are not considered. For example, suppose a certain component experiences demands during a 12-month period such that during seven months demand is zero, during two months the demand is one unit and during the remaining three months the demand is two units. The correct median demand for that component is zero. SASSY would ignore the seven months with zero demand and calculate the median from the remaining five months, with the result being two units.

This programming error was noted in a Naval Audit Report in 1996 and was cited as a major contributing factor to over \$200 million in excess repairable parts. The Marine Corps concurred with the recommendation to calculate median demands using all months and stated that a program fix was transmitted to all activities on 1 August 1996 [2]. We observed, however, that the system continues to calculate medians using only non-zero months.

We also used a revised version of the above formula proposed by the LX office in Headquarters, Marine Corps [13]. Analysts there noticed that the current formula has a mathematical inconsistency: The way SASSY calculates them, RR can be higher than MFR, which can cause RR-MFR to be negative. Also, just setting the safety stock equal to the pipeline stock is inappropriate because the purpose of safety stock is to hedge against variability in the system. So, the LX analysts introduced a variable safety level based on a desired stockout probability. The revised formula is:

$$RO = [p*OST + (1 - p)*RCT]*MFR/30 + SL \quad (3)$$

Where p is the probability that a part is failed beyond repair. We chose to include this formula in our analysis because it illustrates the effect of improving the stock calculation but keeping the current management process.

3. Depot

The Marine Corps Logistics Base (the depot-level inventory and repair facility) does have a program for calculating stock levels. It involves using a forecast of future demand, which is not specified in the documentation. This forecast is then multiplied by the sum of the procurement lead-time and depot repair time. Apparently, however, this formula is not the primary method to determine stock levels. Each end item is managed by an Item Manager who uses experience and judgment to arrive at a level of stock to hold for each repair part. [14]

The VARI-METRIC model calculates total inventory investment for the depot as well as for the RIPs. So, to arrive at a comparable measure, we use a snapshot of stock levels on hand at the depot for calculation of the current total inventory investment. The parts availability experienced at the RIPs is affected by the amounts of stock held at the Depot. It is almost impossible, however, to capture the interaction between the inventory levels using the current information systems. We assume that the contributions of depot stock are subsumed into the order and ship times from the RIPs, so the availability percentages are approximately correct without calculating depot contribution more explicitly. In other words, the OSTs used as inputs to the RIP stock level formula are affected by the depot stock level. If the depot stocks a low amount, then backorders are more common and depot repair times and procurement lead times are added into the OST

quantity. If the depot stocks a large amount, then backorders at the depot are uncommon, and the OST quantity includes only the shipping time between the RIP and the depot.

4. The Model

One of the measures of interest is readily available from the formulas above. Simply multiply each individual RO by the unit price for that part and sum over all parts and all locations to get the total inventory investment. To arrive at an expected availability measure, we note that the current formula models the supply and the maintenance pipelines as M/M/ ∞ queues with OST as the mean service time of the supply queue and RCT as the mean service time of the repair queue. The MFR then, is the arrival rate, but the way it is implemented in this formula is using RR as the probability of successful repair, denoted $(1 - p)$, multiplied by MFR. So if $RR = (1 - p) * MFR$, then $MFR - RR = p * MFR$. The RO, then, is simply the expected number of items in the supply and maintenance pipelines plus a safety level. The full derivation is as follows:

Assume an infinite population of operating systems, each with exponential time between failures. This produces a constant arrival rate, λ (or MFR from above) of failed parts. The probability of washout is still p . If an item does not wash out, it goes to the repair pipeline where it spends an exponentially distributed time, with mean $1/\mu_s$ (or OST), without regard to the number of items already in repair. If an item is not repairable, a replacement is ordered, taking an exponentially distributed time to arrive, with mean $1/\mu_m$ (or RCT), without regard to the number of already outstanding orders. If we define X_s as the number of items on order, X_m as the number of items in maintenance, and X as the total number in either the supply or maintenance pipelines, then $X = X_s + X_m$. Under the given assumptions, both the supply pipeline and the maintenance pipeline are birth-death processes modeled as M/M/ ∞ queues. Therefore, $X_s \sim \text{Poisson}\left(\frac{p\lambda}{\mu_s}\right)$,

$X_m \sim \text{Poisson}\left(\frac{(1-p)\lambda}{\mu_s}\right)$, and $X \sim \text{Poisson}\left(\frac{p\lambda}{\mu_s} + \frac{(1-p)\lambda}{\mu_m}\right)$. Note that the mean of X is equivalent to the expected pipeline stock in the current formula, equation (1). [15]

Using this model, we can calculate the expected number of backorders as a function of the stock level, denoted $EBO(RO)$. If we assume that the inventory is operated under an $(s-1, s)$ policy, a stock-out will occur whenever $X > RO$. So,

$$\begin{aligned} EBO(RO) &= P\{X = RO + 1\} + 2P\{X = RO + 2\} + 3P\{X = RO + 3\} + \dots \\ &= \sum_{x=RO+1}^{\infty} (x - RO)P\{X = x\} \end{aligned} \quad (4)$$

The expected availability of an end item using these parts, can then be calculated using the following equation:

$$A = 100 \prod_{i=1}^I \left[1 - \frac{EBO_i(RO_i)}{NZ_i} \right]^{Z_i} \quad (5)$$

where $i = \{1, \dots, I\}$ are the parts, N is the total number of vehicles, and Z_i is the quantity of part i on each vehicle [7]. This equation assumes that the repair activity does not perform cannibalization, i.e. consolidation of backorders onto the smallest number of vehicles. The availability spoken of here is supply availability, equal to $100[MTBM/(MTBM + MSD)]$, where $MTBM$ is mean time between maintenance and MSD is mean supply delay. For a single vehicle, supply availability is defined as the proportion of time it is not down awaiting parts. For a pool of vehicles it is the proportion of those vehicles that are not down because of parts on backorder.

B. SYSTEM-BASED METHOD

The basic idea of the METRIC model is to minimize the sum of expected backorders across all parts subject to a budget constraint. It performs this minimization by marginal analysis across all parts, all indentures (levels in the part structure, e.g. a valve is part of the head assembly which is part of the engine) and all echelons (levels of repair and inventory activity). METRIC then uses the fact that minimizing backorders is equivalent to maximizing supply availability. The key assumptions are as follows:

- The decision as to whether a base repairs an item does not depend on stock levels or workload. If the base has the capability, it will accomplish the repair regardless of the maintenance workload. If the necessary parts are not available at the base, the base requisitions them from the depot.
- The base is resupplied from the depot, not by lateral redistribution from another base. Despite the recent interest in redistribution, other bases are not a regular

source for parts requisitions. In his thesis on lateral redistribution of repairable parts, Paige [16] found that it is usually not economically beneficial to source parts by redistribution.

- The $(s-1, s)$ inventory policy is appropriate for every item at every echelon. This means that parts are not batched for repair, and that any items beyond repair are reordered on a one-for-one basis.
- Optimal steady state stock levels are determined. This assumes that over some period of time in the future, the number of vehicles or operating hours will remain fairly constant. [7]

We will first describe how the basic model works, and then we will add embellishments. The description of METRIC is summarized from [7].

1. Single-Site Model

First, a single-site model can be derived that computes an optimal curve relating inventory investment to system backorders. A fundamental part of repairable item inventory theory is Palm's Theorem, which enables us to estimate the steady state probability distribution of the number of parts in repair from the probability distribution of demand and the mean of the repair time distribution. It states:

If demand for an item is a Poisson process with annual mean m and if the repair time for each failed unit is independently and identically distributed according to any distribution with mean T years, then the steady-state probability distribution for the number of units in repair has a Poisson distribution with mean mT [7, p. 21].

Sherbrooke defends the assumption of independent repair times as a reasonable approximation by stating that, when queuing does take place, this model will understate repair times, but when management expedites repair, this will overstate repair times and that these two errors will tend to offset each other. The next fundamental basis of the METRIC theory is the equation, $s = OH + DI - BO$. This states that the stock level (s) is always equal to stock on hand (OH) plus stock due in from repair or resupply (DI) minus the number of backorders (BO). Whenever a change occurs in one of the variables on the right-hand side, it is accompanied by a simultaneous change in another. This is also called an $(s-1, s)$ inventory policy, which is reasonable for repairable items because they tend to be high cost and low demand. The model will also use equation (4) for expected

backorders (with the average number of units in repair, mT , from Palm's Theorem as the mean) and equation (5) for availability.

From the mean annual demand and the average repair time we can calculate the expected number of backorders for each part. We use marginal analysis to determine the optimal curve relating system backorders (the sum of part backorders) to inventory cost. To show that marginal analysis produces an optimal solution, Sherbrooke proves that the expected backorder function is convex. Since the expected backorder function is convex, the marginal analysis values $[EBO(s-1) - EBO(s)]/c$, where c denotes the unit cost of an item, are non-increasing. The system backorders are convex also because the sum of convex functions is convex. This marginal analysis procedure will find all of the efficient solutions on the convex hull.

We can show that minimizing the sum of expected backorders is equal to maximizing the availability. For the single-site model, we use equation (5) to calculate the availability for the pool of vehicles. Since the logarithm of a product is equal to the sum of the logarithms, we arrive at:

$$\log(A/100) = \sum_{i=1}^I Z_i \log[1 - EBO_i(s_i)/(NZ_i)] \cong -\sum_{i=1}^I EBO_i(s_i)/N \quad (6)$$

where the last approximation is derived from the power series expansion of $\log(1-a) = -a + 0.5a^2 + \dots$, discarding terms of $O(a^2)$ and higher since a will be small in all cases in which we have interest. Therefore the logarithm of availability is a convex, additive separable function of the item backorder functions. Since a function and its logarithm




Figure 2. Sample Availability vs. Cost Curve

achieve their maximum at the same point, the logarithm of availability is an appropriate surrogate. The desired end result of this marginal analysis procedure is a curve similar to that in Figure 2.

2. Multi-Echelon Model

To expand the single-site model to multiple sites across multiple echelons, we must change the expression for average pipeline stock to account for the interaction between the echelons. We define the following variables for a single part:

- m_j = average annual demand at base j ,
- T_j = average repair time (in years) at base j ,
- μ_j = average pipeline stock at base j ,
- p_j = probability of washout at base j , and
- O_j = average order-and-ship time from depot to base j .

We use the convention that 0 subscripts refer to the depot and positive subscripts refer to the bases. The average annual demand at the depot is the fraction of demand that is not repairable at each base, summed over all bases:

$$m_0 = \sum_{j=1}^J m_j p_j . \quad (7)$$

Since the base demands are assumed to be from Poisson processes, and the sum of Poisson processes is a Poisson process, the base demand is also a Poisson process. The average depot pipeline stock is then $m_0 T_0$, and the expected backorders at the depot, $EBO(s_0 | m_0 T_0)$, are the expected number of supply orders from the bases that are outstanding at the depot at a random point in time. Therefore, the average pipeline stock at a base j is given by:

$$\mu_j = m_j \{ (1 - p_j) T_j + p_j [O_j + EBO(s_0 | m_0 T_0) / m_0] \} \quad j \geq 1 . \quad (8)$$

This average is used as the mean for calculating the expected backorders, as in the single site model. Since the sum of backorders across bases is an additive separable function, we can deal with one item at a time.

The algorithm for multi-echelon optimization is as follows:

1. Start with a depot stock level of zero, $s_0 = 0$.

2. Compute average pipeline stock at each base using equation (8).
3. Calculate the expected backorders for each level of base stock. Repeat for each base.
4. Use marginal analysis to combine the base backorder functions and obtain the minimum backorders for each number of units at the bases. (e.g. For one unit, determine which base to put it such that the sum of base backorders is lowest. For two units, either put both at one base or each at different bases such that the sum of base backorders is minimized, and so forth)
5. If the depot stock level is high enough, go to step 6; otherwise, increase the depot stock level by one and go to step 2.
6. Find the minimum number of expected backorders for each value of total units of stock.
7. Repeat steps 1-6 for each item.
8. Use marginal analysis to combine item solutions, where the first differences must be divided by the item costs. [7]

As we add echelons to the model, we are no longer assured of convexity in the expected backorder function. The non-convex points can be easily identified and can be dealt with by excluding them as potential solutions. We know that the remaining efficient points are convex, so we can use marginal analysis to find the optimum solution.

Since we are now maximizing availability across bases, we must modify the availability function. We take equation (5) and add a subscript, j , for the base to which the availability refers. The combination of base availabilities to compute overall system availability is:

$$A = \sum_{j=1}^J A_j N_j \bigg/ \sum_{j=1}^J N_j \quad (9)$$

This objective function is now the sum of several product functions, so we cannot take logarithms and have it break up into separable pieces. We can, however, get a good approximation to the optimal allocation by finding the maximum across bases and items of the increase in availability times the number of end items divided by the cost of the item – marginal analysis again. Through this procedure, the availability at each base will differ somewhat from one base to another. It is possible, however, to achieve almost equal availabilities at the bases by using a weighting scheme in equation (9). More generally, we can approach any desired set of base availabilities using a similar idea.

3. Modifications and Generalizations

The METRIC model is only an approximation of the optimal solution and tends to understate base backorders, especially in multi-indenture applications. Slay [10] published an improvement to METRIC called VARI-METRIC, which takes into account the variance of the quantity of stock in the pipeline as well as the mean. Graves [11] showed that in 11% of cases, the METRIC stock levels differ by at least one unit from the optimal solution, but the VARI-METRIC levels differ in only 1% of cases. From this, several other modifications were produced that relax restrictive assumptions.

Parts that are common to more than one lower-indenture item, or end-item, were ignored in the previous discussion. They can be handled by apportioning the delay caused by backorders of common items to the lower-indenture parts or end-items according to a binomial probability distribution. Part essentiality differences can be handled by weights applied, by part, to the backorders in a similar manner to how they were applied to base availabilities previously. Other generalizations, such as of the Poisson demand assumption and to allow availability degradation due to maintenance down-time are provided in VARI-METRIC, but we did not use them in order to compare the models on relatively equal footing. We can also relax the no cannibalization and no lateral resupply assumptions but did not in this research, again to compare the models on common ground. We also did not use the multi-indenture capability of the VARI-METRIC model because of the lack of available indenture data for Marine Corps parts. Assuming all repairable components are first indenture will tend to understate availability. Backorders of higher indenture parts will not directly affect the availability of the end-item unless they cause backorders in first indenture parts.

III. DATA AND ASSUMPTIONS

The data for this research were pieced together, sometimes manually, from at least six different sources. Some of the biggest problems the Marine Corps faces in moving to a consolidated inventory or readiness-based sparing are the collection, storage, and retrieval of the required information. Penrose [17] addressed this issue, noting that the Marine Corps could implement a limited RBS capability with the data from current information systems, but would need to make significant changes to realize the full benefits of RBS.

A. CURRENT MODEL

The pieces of data we required for the current and revised models were:

- List of repairable parts for all LAV variants,
- Demand, OST, RCT and washout rates for all parts,
- Numbers of vehicles supported by each inventory site (RIP), and
- Depot stock levels.

We were able to download a list of all parts to each variant from the Item Applications website managed by MCLB. The source, maintenance and recoverability code (SMRC) includes a field that indicates whether a part is consumable or repairable, and if repairable, at which level it is to be repaired. We used the SMRC to segregate the repairable parts. Unfortunately, when a part was common to several variants, its SMRC was not always the same among the various part lists. We used a majority rule to whittle down the number of exceptions. Then, for the few parts for which majority rule did not work, we looked at the nomenclature of the part and used judgment to decide whether or not to include the part in our study. Some parts had incomplete information. We also excluded those parts whose unit price or peacetime replacement factor (PTRF) was absent or equal to zero.

We collected demands, OST, and RCT from three years of data (fiscal years 1998 – 2000) extracted from MIMMS by the Secondary Repairables Management Office at Marine Corps Logistics Base, Albany. This data indicated the month in which a demand

occurred. We could, therefore, aggregate the demands by month and calculate the median and average monthly demands independently of the Marine Corps supply system, SASSY. Recall that SASSY calculates medians and averages using only the months with non-zero demand. We calculated medians in the same way so that we would produce equivalent ROs. We calculated true average demands, however, using all months in the observed period.

We noticed from these data that only about 15% of the CEC-5 parts had recorded demands in the three-year period covered by our data. We considered calculating an estimated demand for those that did not have a recorded demand using the PTRF. The PTRF is defined as “the average rate at which the type of item has been used by Marine Corps field units or the rate at which design engineers anticipate the item will fail, wear out, or otherwise require replacement” [18]. It refers to the average proportion of items that are expected to need replacement in a year. We discovered, however, that several of the values were obviously in error. For example, several items, such as the engine and fuel tank, had a PTRF equal to 1 meaning that, on average, a given part would fail once each year. These components had no registered demand in during the period considered. Because of these errors and not knowing how widespread they might be, we chose not to use the PTRF to estimate any demands.

Leaving the average monthly demand equal to zero for those parts with no demands during the three year period is somewhat unsatisfying because there must be some level of expected demand, though we do know it must be relatively small. To explore the effect of these treatments of zero demand items, we considered three cases: leaving the parts with zero demand as zeros, inserting a quantity equal to one demand in 10 years, and inserting a quantity equal to one demand in five years.

The demand data included a code that indicated if a failed part was successfully repaired. We estimated the washout rate, p , as unsuccessful repairs divided by total demands. If a part had no recorded demand, we set washout rate at the RIP equal to one for all depot-level repairables (DLR), as denoted by a SMRC ending in “DD”. This is because the definition of a depot-level repairable is a part that can only, at least according

to policy, be repaired at the depot. For all field-level repairables (FLR), we assumed an arbitrary washout rate at the RIP equal to 0.2.

The Depot apparently did not have stock levels to which they firmly held. We used a list of the stock held at the both MCLB Albany and Barstow as of 25 July 2001. The inventory is separated by codes according to the purpose for which it is held and the condition it is in. We included only those items that were in operating stock as indicated by purpose code. Further, we included those components that were serviceable, unserviceable (but repairable), or currently undergoing maintenance as indicated by the condition code.

For the quantities of vehicles supported by each RIP, we use the numbers from the official tables of equipment. The inventory in I MEF supports First and Third Light Armored Reconnaissance (LAR) Battalions, the Equipment Allowance Pool and the School of Infantry. The inventory in II MEF supports Second LAR Battalion and a few additional assets at Aberdeen Proving Ground. III MEF supports the LAV company in the Combat Assault Battalion. We assume that one company from Third LAR Battalion is deployed from I MEF to III MEF.

B. DATA FOR VMETRIC™-XL

The data requirements for running the METRIC-based model were driven by the software we chose to use, VMetric™-XL. As much as possible we used the same data as with the current and revised formulas. Where different information was required we derived it from the original data.

Demand in VMetric is input as either MTBF or MRR6, where MRR6 is defined as the number of failures per million operating hours. Since we already had the number of failures per month, we only needed an average of operating hours per month to calculate MRR6. We received the average annual operating hours per vehicle for each year from 1992 to 1997 for each LAV variant from the LAV Program Office. To estimate usage, we used an exponentially weighted average, with a weight of 0.7, of the yearly figures in which the more recent years were weighted more heavily. We calculated MRR6 for each part at each RIP as follows: The annual demand produced by the pool of parts supported by a particular RIP is simply the monthly demand multiplied

by 12 months. The total annual hours that same quantity of parts was used is the average annual operating hours per vehicle multiplied the quantity of vehicles that use that part multiplied by the number of that part used in each vehicle. MRR6 is total annual demand divided by annual operating hours multiplied by one million.

We used the same values of washout rate, specific to each item at each base, in VMetric as in the current and revised models. We used specific RCT values wherever available, and a default value at each base equal to the average repair time at that base. The OST in VMetric was defined a little differently, so we had to make some adjustments. We needed a value for the OST between each base and the depot, given that the requested item was on the shelf at the depot. The OST value from the Marine Corps supply system is the total time from when a RIP orders a part from the depot until it is received, including any delays for repair or resupply at the depot. After looking at the distribution of OST times for each RIP, we chose a value of 15 days for all bases. For each specific part we also needed a procurement lead-time (PLT), which is the same thing as OST between the depot and the vendor. We set this value to the OST from the RIP recomputation minus 15 days. This way, the total ordering delay is approximately equal in both cases.

VMetric also required information about the depot. After discussions with the Secondary Repairables Management Office (SRMO) at MCLB, we used a standard rate of 60 days for RCT at the depot. The SRMO also stated that most things are successfully repaired. Therefore, we used a washout rate of 10% for FLRs and 20% for DLRs. The rationale for the difference was that FLRs are so designated because they are generally simpler to repair.

IV. RESULTS AND ANALYSIS

The current model and the stock levels it computes are the baseline against which all others will be compared. The total investment in mission essential repairable parts for the LAV is \$25.9 million, giving a computed supply availability of 89.1%. We show that just changing the method of calculating median demand will drastically reduce investment, but may reduce availability beyond what is acceptable. The revised formula provided by LX produces marginally improved results in both total investment and availability. The VMetric model, however, results in an almost four-fold reduction in inventory investment while increasing availability.

A. CHANGES WITHIN THE CURRENT MODEL

Before going into the full comparison of the results of the current model and the centralized model, we observe the results from the current model and look at the effect of the way SASSY computes median demands. We note in Table 1, that both II and III MEFs hold a higher dollar value of parts, despite operating many fewer vehicles than I MEF. For a reason we could not determine, II MEF experiences the same or higher level of demand as I MEF. In addition, it appears that II MEF had a generally longer RCT than the other MEFs and experienced unusually high demands for a few expensive parts. This seems to explain the larger inventory investment. Despite the greater quantity of stock, the availability is lower than I MEF because of a smaller quantity of LAVs. In other words, if I MEF and II MEF have the same number of expected backorders, II MEF's availability will be lower because those backorders are spread over fewer vehicles. III MEF also seemed to experience a higher rate of demand than the small number of vehicles would seem to justify. We observed for III MEF that a few expensive parts experienced high washout rates and long OST. These few items were enough to drive the inventory cost above that of I MEF. We also note that III MEF stocked only 65 line items compared to 98 for the other two sites. This seems to give a reason for the lower availability at III MEF.

The total inventory costs and availabilities calculated here might not look the same as those experienced in these organizations. Inventory managers at the RIPs

generally do not trust the stock levels provided by SASSY and use their judgment to modify those quantities. The modifications tend to reduce the overall inventory investment [19]; but, due to the lack of these modified stock levels, we are unable to determine their effect on availability.

	# of LAVs	Current Method		True Median		Mean	
		Invest.	Avail.	Invest.	Avail.	Invest.	Avail.
I MEF	241	\$3.12M	92.0%	\$1.26M	69.1%	\$2.03M	89.6%
II MEF	122	\$5.78M	88.4%	\$1.39M	14.9%	\$3.99M	79.0%
III MEF	50	\$3.84M	76.9%	\$0.15M	19.3%	\$0.49M	56.5%
Depot		\$13.18M		\$13.18M		\$13.18M	
Total	413	\$25.92M	89.1%	\$15.99M	47.0%	\$19.69M	82.5%

Table 1. Effect of Various Median Calculations

As shown in Table 1, if medians were calculated correctly, by including all monthly observations, the total parts investment would drop more than \$10 million, a reduction of 38%. One problem with this is that it may reduce stock levels too much. The computed supply availability for the LAV using true medians is 47%. The reason for this drop is that repairables tend to experience low demand. As such, it is common to see a median demand of zero even when the annual demand is significant. For I MEF, 90% of the true median demands were less than the current, inflated, median. For II MEF and III MEF, over 95% of the true medians were smaller. To put it another way, using true medians, at I MEF only ten line items had non-zero demand, meaning those items received stock; but, at II MEF and III MEF, only four parts had non-zero demand, so only those four received stock. The low availability levels indicate that we are on the steep part of the availability curve, which means any small change in inventory cost translates to a large change in availability. This could explain the dramatic difference in the effect of the different median calculation on availability.

Another alternative to medians is to simply use the sample mean. To calculate the numbers shown in Table 1, we used sample means rather than medians for all quantities in the current formula. This adjustment realizes a 24% reduction in inventory for a relatively small reduction in availability. The formula using means reduces stock level

from the current method because we used all monthly observations, including zeros, in the calculation. It also makes sense that total investment is higher using means than true medians because the demand distributions tend to have a heavy right tail.

The results from the revised formula are shown in Table 2. It does reduce overall inventory cost because it includes a variable safety level. In some cases, not much safety level is needed because of low variability in demand, OST and RCT. In other cases, the revised formula increases the stock quantity from the baseline because of high variability. It also increases availability because it ensures that each part is stocked in sufficient

	Current Formula		Revised Formula	
	Investment	Availability	Investment	Availability
I MEF	\$3.12M	92.0%	\$2.51M	97.7%
II MEF	\$5.78M	88.4%	\$4.26M	92.9%
III MEF	\$3.84M	76.9%	\$1.89M	91.9%
Depot	\$13.18M		\$13.18M	
Total	\$25.92M	89.1%	\$21.84M	95.6%

Table 2. Comparison of Revised and Current Formulas.

quantity to be at least 95% confident that no stock-out will occur. These effects aggregate to a \$4.4 million reduction in inventory and an increase of 6.4% in availability. The LX analysts note that the improvements gained by using this revised formula are almost identical to those from the inventory managers at each RIP using experience and judgment to modify the ROs produced by SASSY [19].

B. CHANGING THE MODEL

The previous results come only from minor changes in formulas within the existing model. If we change the whole inventory model, we see much more dramatic results. With VMetric any level of availability can be achieved as long as enough is invested. At the same level of availability as the current model (89.1%), VMetric recommends total stock levels worth \$2.9 million, a reduction in inventory investment of 89%. To achieve 99% availability, VMetric recommends an investment of \$6.9 million. A chart displaying these results is shown in Figure 3.




Figure 3. Optimal Availability vs. Cost Curve for LAV Repairables.

The curve represents the maximum availability achieved at each level of investment. An additional output are the optimal stock levels by part and site that achieve that availability. Discussion with executives at Systems Exchange, the company that developed the VMetric-XL software, revealed that two-fold reductions in inventory investment are common in their experience. They said further that they often expect three-fold improvement in Department of Defense implementations, and that four-fold improvement is within reason [20, 21].

VMetric can achieve such reductions for a number of reasons: It captures the interaction between the RIPs and the Depot much more explicitly than the current model with computed stock levels and expected backorders given the stock levels at the depot. Also, it is an optimization model – maximizing availability given investment. We observe that the VARI-METRIC algorithm places the majority of the stock at the RIPs and very little at the Depot. A summary by organization for 99% availability is shown in Table 3. The reason for such a dramatic drop in stock at the Depot is that VMetric will tend to place stock at the RIPs because of the relatively long OST assumed between the RIPs and the Depot. The VARI-METRIC model can reduce the inventory investment at

the RIPs because of the correction in calculation of the average demand, and because it chooses the optimal mix of parts and places them optimally among the inventory sites.

	Current Method		VMetric	
	Investment	Availability	Investment	Availability
I MEF	\$3.12M	92.0%	\$2.14M	99.5%
II MEF	\$5.78M	88.4%	\$3.80M	98.5%
III MEF	\$3.84M	76.9%	\$0.82M	98.0%
Depot	\$13.18M		\$0.15M	
Total	\$25.92M	89.1%	\$6.91M	99.0%

Table 3. Direct Comparison of Current Model vs. VMetric

The fact that each of the models above, as implemented, ignores indenture information and treats all parts as first indenture level will tend to underestimate availability. On the other hand, both models assume demands to come from a Poisson distribution, which will tend to overstate availability, because the variance of a Poisson distribution is equal to the mean and the actual variance of most parts in this study is probably greater than the mean. Since we underestimate the variance, the availability achieved from a given stock level will appear higher than the true value. We cannot, however, quantify these effects within the scope of this research.

As mentioned previously there was some question about what to do with the parts that had no recorded demand during the three-year observation period. The expected demand must be some positive number, however small, because every part must fail at some time. Also, leaving the demand equal to zero would cause VMetric to always set the stock levels for those parts equal to zero. The results mentioned above all leave the zero demands as true zeros. We also ran the models with a quantity equal to one demand in 10 years, and again with one demand in five years, inserted in the places where expected demand was zero. The effects of these changes on the current model are shown in Table 4. These changes do not affect the investment quantity because we are only changing the expected demand of those items that SASSY did not stock. Because we do not change any stock levels there are now backorders expected for each of these parts, so

	Investment	Availability		
		True Zeros	1 Demand in 10 Years	1 Demand in 5 Years
I MEF	\$3.12M	92.0%	79.9%	69.4%
II MEF	\$5.78M	88.4%	74.0%	61.9%
III MEF	\$3.84M	76.9%	67.1%	58.0%
Depot	\$13.18M			
Total	\$25.92M	89.1%	76.6%	65.8%

Table 4. Modifying Expected Demands in the Current Model

there is a significant reduction in availability. The expected backorder quantities are small, so the reduction in availability is small for each individual part, but still makes the availability of every part less than one. The total availability for the system is a multiplicative function, and the result of any number less than one multiplied by itself hundreds of times over is a very small number. The effects of modifying the items with zero expected demand in this way are much less in VMetric because the VARI-METRIC algorithm starts from scratch and builds the optimal parts mix in each case. At an investment of \$2.9M with true zeros, the availability is 89.2%, with one demand in 10 years it is 88.8%, and with one demand in five years it is 87.9%.

C. CONCLUSIONS

Our results show that improvements can be achieved by affecting small changes to the calculation methods within the current inventory model. The current method of computing median demands and repair rates is erroneous and contributes to excess inventory. Fixing the median calculation method will reduce inventory but it also may decrease availability because such a correction still does not address the underlying inventory model. Slight changes to the formula can likewise produce improvements but, again, do not affect the inventory model.

Most importantly, the results show that there are huge potential benefits from managing the Marine Corps inventory as a whole rather than as a disjoint group of independent inventories. Centralized management of secondary repairables cannot mean simply consolidating ownership under one command or funding all RIPs through a single

command. To get the maximum benefits, Marine Corps secondary repairables must be treated as one inventory separated among geographically separate sites, not as separate inventories. Using the same modeling assumptions, we expect that similar results could be achieved with any other item of equipment in the Marine Corps inventory. If the Marine Corps secondary repairables were centrally managed using a readiness-based sparing model based on the METRIC or VARI-METRIC algorithms, such as VMetric™-XL, we would expect large cost savings. However, we expect that improvements from implementing a tool such as VMetric would not be as large as those indicated in our results because of the stock level modifications performed by inventory managers under the current system.

Additionally, we note that an inventory model such as VARI-METRIC is more data intensive than the current model. As discussed in [17], the Marine Corps could implement a readiness-based sparing model using data elements currently captured, but could generate a much more accurate solution from such a model using additional information. If implemented, we recommend adding indenture and variance of demand information, a more accurate breakdown of inventory sources and sites, more accurate procurement lead times, and a more thorough analysis of expected demands. From our experiences in this research we add that the quality of existing data, such as SMRC, PTRF and operational usage, needs to be improved. As noted in [22], the quality of the solution from a readiness-based sparing model is dependent on the accuracy of the data, and how current they are.

THIS PAGE INTENTIONALLY LEFT BLANK

LIST OF REFERENCES

1. Interviews between logistics analysts at Headquarters, Marine Corps (LX) and the author during September 2000.
2. Naval Audit Service, *Inventory Requirements at Marine Corps Repairable Issue Points (059-96)*. Naval Audit Report, 13 September 1996.
3. Commandant of the Marine Corps White Letter of Jan 97.
4. Commandant of the Marine Corps (LP) message, CMC(LP) 010101Z MAY 00.
5. Commander, Marine Corps Materiel Command message, COMMARCORMATCOM 260100Z JUN 00.
6. Interview between Larry Paige and the author on 5 December 2000.
7. Craig C. Sherbrooke, *Optimal Inventory Modeling of Systems: Multi-Echelon Techniques*, John Wiley and Sons, Inc., 1992.
8. Angel Diaz and Michael C. Fu, *Models for multi-echelon repairable item inventory systems with limited repair capacity*, European Journal of Operational Research 97, (1997) pp. 480-492.
9. Yunzeng Wang, Morris A. Cohen and Yu-Sheng Zheng, *A Two-Echelon Repairable Inventory System with Stocking-Center-Dependent Depot Replenishment Lead Times*, Management Science 46, pp. 1441-1453.
10. F. M. Slay, *VARI-METRIC: An Approach to Modeling Multi-Echelon Resupply when the Demand Process is Poisson with a Gamma Prior*. Logistics Management Institute, Washington, D.C. Report AF301-3.
11. S. C. Graves, *A Multi-Echelon Inventory Model for a Repairable Item with One for-One Replenishment*, Management Science 31, pp. 1247-1256.
12. United States Marine Corps, *MCO P4400.151B with Ch. 3: FMF SASSY Management Unit Procedures*, 28 April 1989.
13. Headquarters, Marine Corps (LX) presentation, *Assessing the RIP RO Recomputation Formula*, 15 August 2000.
14. Larry Paige, electronic mail message, 25 July 2001.

15. Chris Goodheart, handwritten notes.
16. Larry G. Paige II, *Inventory Management of Repairables in the U. S. Marine Corps – A Virtual Warehouse Concept*, thesis, Naval Postgraduate School, June 2000.
17. Craig B. Penrose, *Data Requirements for Availability Based Sparing in the U. S. Marine Corps*, thesis, Naval Postgraduate School, September 1998.
18. SL-4-08757A, *Marine Corps Stocklist: Tractor, Full-Track*, April 1986.
19. Headquarters, Marine Corps (LX) presentation, *Assessing the RIP RO Recomputation Formula, Part III*, 15 May 2001
20. Bob Butler, electronic mail message to the author, 5 April 2001.
21. Interview between Chris Hampson and the author on 23 August 2001.
22. Anne J. Hale, *Analysis of America's Readiness-Based Sparing Aviation Consolidated Allowance List (CRM 94-140)*, Center for Naval Analyses, December 1994.

INITIAL DISTRIBUTION LIST

1. Defense Technical Information Center
Ft. Belvoir, Virginia
2. Dudley Knox Library
Naval Postgraduate School
Monterey, California
3. Marine Corps Representative
Naval Postgraduate School
Monterey, California

4. Director, Training and Education, MCCDC, Code C46
Quantico, Virginia

5. Director, Marine Corps Research Center, MCCDC, Code C40RC
Quantico, Virginia

6. Marine Corps Tactical Systems Support Activity (Attn: Operations Officer)
Camp Pendleton, California

7. Director, Studies and Analysis Division, MCCDC, Code C45
Quantico, Virginia

8. Professor Kevin Gue
Naval Postgraduate School
Monterey, California

9. Professor David Schrady
Naval Postgraduate School
Monterey, California

10. Lieutenant Colonel Greg Mislick
Naval Postgraduate School
Monterey, California

11. Secondary Repairables Management Office
Marine Corps Logistics Base
Albany, Georgia

12. Headquarters, Marine Corps (LX)
Arlington, Virginia

13. Deputy Chief of Naval Operations (Logistics)
Attn: LCDR Arthur Cimiluca, N421C
Washington, District of Columbia

14. Systems Exchange
Pacific Grove, California
